

UNITED STATES LETTERS PATENT APPLICATION
FOR:

**METHOD AND SYSTEM FOR PROCESSING QUERY MESSAGES
OVER A NETWORK**

INVENTOR:

Aristotle N. Balogh

ASSIGNEE:

VeriSign, Inc.

Prepared by:

KENYON & KENYON
1500 K Street, NW, Suite 700
Washington, DC 20005
(202) 220-4200

METHOD AND SYSTEM FOR PROCESSING QUERY MESSAGES OVER A NETWORK

Claim For Priority/Cross Reference to Related Applications

[0001] This non-provisional application claims the benefit of U.S. Provisional Patent Application No. 60/330,842, filed November 1, 2001, which is incorporated by reference in its entirety, and U.S. Provisional Patent Application No. 60/365,169, filed March 19, 2002, which is incorporated by reference in its entirety. This application is related to U.S. Non-Provisional Patent Application Serial Nos. [Att'y Dkt 12307/100178], [Att'y Dkt 12307/100180], [Att'y Dkt 12307/100181] and [Att'y Dkt 12307/100182].

Technical Field

[0002] The present invention relates to computer systems. More specifically, the present invention relates to a method and system for processing query messages over a network.

Background of the Invention

[0003] As the Internet continues its meteoric growth, scaling domain name service (DNS) resolution for root and generic top level domain (gTLD) servers at reasonable price points is becoming increasingly difficult. The A root server (i.e., a.root-server.net) maintains and distributes the Internet namespace root zone file to the 12 secondary root servers geographically distributed around the world (i.e., b.root-server.net, c.root-server.net, etc.), while the corresponding gTLD servers (i.e., a.gtld-servers.net, b.gtld-servers.net, etc.) are similarly distributed and support the top level domains (e.g., *.com, *.net, *.org, etc.). The ever-increasing volume of data coupled with the unrelenting growth in query rates is forcing a complete rethinking of the hardware and software infrastructure needed for root and gTLD DNS service over the next several years. The typical single server installation of the standard "bind" software distribution is already insufficient for the demands of the A root and will soon be unable to meet even gTLD needs. With the convergence of the public switched telephone network (PSTN) and the Internet, there are opportunities for a general purpose, high performance search mechanism to provide features normally associated with Service Control Points (SCPs) on the PSTN's SS7 signaling network as new, advanced services are offered that span

the PSTN and the Internet, including Advanced Intelligent Network (AIN), Voice Over Internet Protocol (VoIP) services, geolocation services, etc.

Brief Description of the Drawings

[0004] FIG. 1 is a system block diagram, according to an embodiment of the present invention.

[0005] FIG. 2 is a detailed block diagram that illustrates a message data structure, according to an embodiment of the present invention.

[0006] FIG. 3 is a detailed block diagram that illustrates a message latency data structure architecture, according to an embodiment of the present invention.

[0007] FIG. 4 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention.

[0008] FIG. 5 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention.

[0009] FIG. 6 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention.

[0010] FIG. 7 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention.

[0011] FIG. 8 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention.

[0012] FIG. 9 is a top level flow diagram that illustrates a method for processing query messages over a network, according to an embodiment of the present invention.

[0013] FIG. 10 is a top level flow diagram that illustrates a method for determining a message latency associated with a sequence number, according to an embodiment of the present invention.

[0014] FIG. 11 is a top level flow diagram that illustrates a method for determining a message latency associated with a sequence number, according to an embodiment of the present invention.

Detailed Description

[0015] Embodiments of the present invention provide a method and system for processing query messages over a network. Specifically, a plurality of queries may be extracted from a plurality of query messages received from a plurality of users over a network. A number of queries, included in the plurality of queries, may be determined, and a current sequence number may be associated with the plurality of queries. A request message may be created including the plurality of queries, a first sequence number equal to the current sequence number and a first message count equal to the number of queries. The request message may be sent to a search engine and a response message may be received from the search engine. The response message may include a plurality of replies, a second sequence number, a second message count, a third sequence number and a third message count. A plurality of reply messages may be created from the plurality of replies and sent to the plurality of users over the network.

[0016] FIG. 1 is a block diagram that illustrates a system according to an embodiment of the present invention. Generally, system 100 may host a large, memory-resident database, receive search requests and provide search responses over a network. For example, system 100 may be a symmetric, multiprocessing (SMP) computer, such as, for example, an IBM RS/6000® M80 or S80 manufactured by International Business Machines Corporation of Armonk, New York, a Sun Enterprise™ 10000 manufactured by Sun Microsystems, Inc. of Santa Clara, California, etc. System 100 may also be a multi-processor personal computer, such as, for example, a Compaq ProLiant™ ML530 (including two Intel Pentium® III 866 MHz processors) manufactured by Hewlett-Packard Company of Palo Alto, California. System 100 may also include a multiprocessing operating system, such as, for example, IBM AIX® 4, Sun Solaris™ 8 Operating Environment, Red Hat Linux® 6.2, etc. System 100 may receive periodic updates over network 124, which may be concurrently incorporated into the database. Embodiments of the present invention may achieve very high database search and update throughput by incorporating each update to the database without the use of database locks or access controls.

[0017] In an embodiment, system 100 may include at least one processor 102-1 coupled to bus 101. Processor 102-1 may include an internal memory cache (e.g., an L1 cache, not shown for clarity). A secondary memory cache 103-1 (e.g., an L2 cache, L2/L3 caches, etc.) may reside between processor 102-1 and bus 101. In a preferred

embodiment, system 100 may include a plurality of processors 102-1 ... 102-P coupled to bus 101. A plurality of secondary memory caches 103-1 ... 103-P may also reside between plurality of processors 102-1 ... 102-P and bus 101 (e.g., a look-through architecture), or, alternatively, at least one secondary memory cache 103-1 may be coupled to bus 101 (e.g., a look-aside architecture). System 100 may include memory 104, such as, for example, random access memory (RAM), etc., coupled to bus 101, for storing information and instructions to be executed by plurality of processors 102-1 ... 102-P.

[0018] Memory 104 may store a large database, for example, for translating Internet domain names into Internet addresses, for translating names or phone numbers into network addresses, for providing and updating subscriber profile data, for providing and updating user presence data, etc. Advantageously, both the size of the database and the number of translations per second may be very large. For example, memory 104 may include at least 64 GB of RAM and may host a 500M (i.e., 500×10^6) record domain name database, a 500M record subscriber database, a 450M record telephone number portability database, etc.

[0019] On an exemplary 64-bit system architecture, such as, for example, a system including at least one 64-bit big-endian processor 102-1 coupled to at least a 64-bit bus 101 and a 64-bit memory 104, an 8-byte pointer value may be written to a memory address on an 8-byte boundary (i.e., a memory address divisible by eight, or, e.g., 8N) using a single, uninterruptible operation. Generally, the presence of secondary memory cache 103-1 may simply delay the 8-byte pointer write to memory 104. For example, in one embodiment, secondary memory cache 103-1 may be a look-through cache operating in write-through mode, so that a single, 8-byte store instruction may move eight bytes of data from processor 102-1 to memory 104, without interruption, and in as few as two system clock cycles. In another embodiment, secondary memory cache 103-1 may be a look-through cache operating in write-back mode, so that the 8-byte pointer may first be written to secondary memory cache 103-1, which may then write the 8-byte pointer to memory 104 at a later time, such as, for example, when the cache line in which the 8-byte pointer is stored is written to memory 104 (i.e., e.g., when the particular cache line, or the entire secondary memory cache, is "flushed").

[0020] Ultimately, from the perspective of processor 102-1, once the data are latched onto the output pins of processor 102-1, all eight bytes of data are written to memory

104 in one contiguous, uninterrupted transfer, which may be delayed by the effects of a secondary memory cache 103-1, if present. From the perspective of processors 102-2 ... 102-P, once the data are latched onto the output pins of processor 102-1, all eight bytes of data are written to memory 104 in one contiguous, uninterrupted transfer, which is enforced by the cache coherency protocol across secondary memory caches 103-1 ... 103-P, which may delay the write to memory 104 if present.

[0021] However, if an 8-byte pointer value is written to a misaligned location in memory 104, such as a memory address that crosses an 8-byte boundary, all eight bytes of data can not be transferred from processor 102-1 using a single, 8-byte store instruction. Instead, processor 102-1 may issue two separate and distinct store instructions. For example, if the memory address begins four bytes before an 8-byte boundary (e.g., $8N - 4$), the first store instruction transfers the four most significant bytes to memory 104 (e.g., $8N - 4$), while the second store instruction transfers the four least significant bytes to memory 104 (e.g., $8N$). Importantly, between these two separate store instructions, processor 102-1 may be interrupted, or, processor 102-1 may lose control of bus 101 to another system component (e.g., processor 102-P, etc.). Consequently, the pointer value residing in memory 104 will be invalid until processor 102-1 can complete the second store instruction. If another component begins a single, uninterruptible memory read to this memory location, an invalid value will be returned as a presumably valid one.

[0022] Similarly, a new 4-byte pointer value may be written to a memory address divisible by four (e.g., $4N$) using a single, uninterruptible operation. Note that in the example discussed above, a 4-byte pointer value may be written to the $8N - 4$ memory location using a single store instruction. Of course, if a 4-byte pointer value is written to a location that crosses a 4-byte boundary, e.g., $4N - 2$, all four bytes of data can not be transferred from processor 102-1 using a single store instruction, and the pointer value residing in memory 104 may be invalid for some period of time.

[0023] System 100 may also include a read only memory (ROM) 106, or other static storage device, coupled to bus 101 for storing static information and instructions for processor 102-1. A storage device 108, such as a magnetic or optical disk, may be coupled to bus 101 for storing information and instructions. System 100 may also include display 110 (e.g., an LCD monitor) and input device 112 (e.g., keyboard, mouse, trackball, etc.), coupled to bus 101. System 100 may include a plurality of network interfaces 114-1 ... 114-O, which may send and receive electrical, electromagnetic or

optical signals that carry digital data streams representing various types of information. In an embodiment, network interface 114-1 may be coupled to bus 101 and local area network (LAN) 122, while network interface 114-O may be coupled to bus 101 and wide area network (WAN) 124. Plurality of network interfaces 114-1 ... 114-O may support various network protocols, including, for example, Gigabit Ethernet (e.g., IEEE Standard 802.3-2002, published 2002), Fiber Channel (e.g., ANSI Standard X.3230-1994, published 1994), etc. Plurality of network computers 120-1 ... 120-N may be coupled to LAN 122 and WAN 124. In one embodiment, LAN 122 and WAN 124 may be physically distinct networks, while in another embodiment, LAN 122 and WAN 124 may be via a network gateway or router (not shown for clarity). Alternatively, LAN 122 and WAN 124 may be the same network.

[0024] As noted above, system 100 may provide DNS resolution services. In a DNS resolution embodiment, DNS resolution services may generally be divided between network transport and data look-up functions. For example, system 100 may be a back-end look-up engine (LUE) optimized for data look-up on large data sets, while plurality of network computers 120-1 ... 120-N may be a plurality of front-end protocol engines (PEs) optimized for network processing and transport. The LUE may be a powerful multiprocessor server that stores the entire DNS record set in memory 104 to facilitate high-speed, high-throughput searching and updating. In an alternative embodiment, DNS resolution services may be provided by a series of powerful multiprocessor servers, or LUEs, each storing a subset of the entire DNS record set in memory to facilitate high-speed, high-throughput searching and updating.

[0025] Conversely, the plurality of PEs may be generic, low profile, PC-based machines, running an efficient multitasking operating system (e.g., Red Hat Linux® 6.2), that minimize the network processing transport load on the LUE in order to maximize the available resources for DNS resolution. The PEs may handle the nuances of wire-line DNS protocol, respond to invalid DNS queries and multiplex valid DNS queries to the LUE over LAN 122. In an alternative embodiment including multiple LUEs storing DNS record subsets, the PEs may determine which LUE should receive each valid DNS query, and multiplex valid DNS queries to the appropriate LUEs. The number of PEs for a single LUE may be determined, for example, by the number of DNS queries to be processed per second and the performance characteristics of the particular system. Other metrics may also be used to determine the appropriate mapping ratios and behaviors.

[0026] Generally, other large-volume, query-based embodiments may be supported, including, for example, telephone number resolution, SS7 signaling processing, geolocation determination, telephone number-to-subscriber mapping, subscriber location and presence determination, etc.

[0027] In an embodiment, a central on-line transaction processing (OLTP) server 140-1 may be coupled to WAN 124 and receive additions, modifications and deletions (i.e., update traffic) to database 142-1 from various sources. OLTP server 140-1 may send updates to system 100, which includes a local copy of database 142-1, over WAN 124. OLTP server 140-1 may be optimized for processing update traffic in various formats and protocols, including, for example, HyperText Transmission Protocol (HTTP), Registry Registrar Protocol (RRP), Extensible Provisioning Protocol (EPP), Service Management System/800 Mechanized Generic Interface (MGI), and other on-line provisioning protocols. A constellation of read-only LUEs may be deployed in a hub and spoke architecture to provide high-speed search capability conjoined with high-volume, incremental updates from OLTP server 140-1.

[0028] In an alternative embodiment, data may be distributed over multiple OLTP servers 140-1...140-S, each of which may be coupled to WAN 124. OLTP servers 140-1...140-S may receive additions, modifications, and deletions (i.e., update traffic) to their respective databases 142-1...142-S (not shown for clarity) from various sources. OLTP servers 140-1...140-S may send updates to system 100, which may include copies of databases 142-1...142-S, other dynamically-created data, etc., over WAN 124. For example, in a geolocation embodiment, OLTP servers 140-1...140-S may receive update traffic from groups of remote sensors. In another alternative embodiment, plurality of network computers 120-1 ... 120-N may also receive additions, modifications, and deletions (i.e., update traffic) from various sources over WAN 124 or LAN 122. In this embodiment, plurality of network computers 120-1 ... 120-N may send updates, as well as queries, to system 100.

[0029] In the DNS resolution embodiment, each PE (e.g., each of the plurality of network computers 120-1 ... 120-N) may combine, or multiplex, several DNS query messages, received over a wide area network (e.g., WAN 124), into a single Request SuperPacket and send the Request SuperPack t to the LUE (e.g., system 100) over a local area network (.g., LAN 122). The LUE may combine, or multiplex, several DNS query message replies into a single Response SuperPacket and send the R sponse

SuperPacket to the appropriate PE over the local area network. Generally, the maximum size of a Request or Response SuperPacket may be limited by the maximum transmission unit (MTU) of the physical network layer (e.g., Gigabit Ethernet). For example, typical DNS query and reply message sizes of less than 100 bytes and 200 bytes, respectively, allow for over 30 queries to be multiplexed into a single Request SuperPacket, as well as over 15 replies to be multiplexed into a single Response SuperPacket. However, a smaller number of queries (e.g., 20 queries) may be included in a single Request SuperPacket in order to avoid MTU overflow on the response (e.g., 10 replies). For larger MTU sizes, the number of multiplexed queries and replies may be increased accordingly.

[0030] Each multitasking PE may include an inbound thread and an outbound thread to manage DNS queries and replies, respectively. For example, the inbound thread may unmarshal the DNS query components from the incoming DNS query packets received over a wide area network and multiplex several milliseconds of queries into a single Request SuperPacket. The inbound thread may then send the Request SuperPacket to the LUE over a local area network. Conversely, the outbound thread may receive the Response SuperPacket from the LUE, de-multiplex the replies contained therein, and marshal the various fields into a valid DNS reply, which may then be transmitted over the wide area network. Generally, as noted above, other large-volume, query-based embodiments may be supported.

[0031] In an embodiment, the Request SuperPacket may also include state information associated with each DNS query, such as, for example, the source address, the protocol type, etc. The LUE may include the state information, and associated DNS replies, within the Response SuperPacket. Each PE may then construct and return valid DNS reply messages using the information transmitted from the LUE. Consequently, each PE may advantageously operate as a stateless machine, i.e., valid DNS replies may be formed from the information contained in the Response SuperPacket. Generally, the LUE may return the Response SuperPacket to the PE from which the incoming SuperPacket originated; however, other variations may obviously be possible.

[0032] In an alternative embodiment, each PE may maintain the state information associated with each DNS query and include a reference, or handle, to the state information within the Request SuperPacket. The LUE may include the state information references, and associated DNS replies, within the Response SuperPacket. Each PE

may then construct and return valid DNS reply messages using the state information references transmitted from the LUE, as well as the state information maintained thereon. In this embodiment, the LUE may return the Response SuperPacket to the PE from which the incoming SuperPacket originated.

[0033] FIG. 2 is a detailed block diagram that illustrates a message data structure, according to an embodiment of the present invention. Generally, message 200 may include header 210, having a plurality of sequence number 211-1 ... 211-S and a plurality of message counts 212-1 ... 212-S, and data payload 215.

[0034] In the DNS resolution embodiment, message 200 may be used for Request SuperPackets and Response SuperPackets. For example, Request SuperPacket 220 may include header 230, having a plurality of sequence number 231-1 ... 231-S and a plurality of message counts 232-1 ... 232-S, and data payload 235 having multiple DNS queries 236-1 ... 236-Q, accumulated by a PE over a predetermined period of time, such as, for example, several milliseconds. In one embodiment, each DNS query 236-1 ... 236-Q may include state information, while in an alternative embodiment, each DNS query 236-1 ... 236-Q may include a handle to state information.

[0035] Similarly, Response SuperPacket 240 may include header 250, having a plurality of sequence number 251-1 ... 251-S and a plurality of message counts 252-1 ... 252-S, and data payload 255 having multiple DNS replies 256-1 ... 256-R approximately corresponding to the multiple DNS queries contained within Request SuperPacket 220. In one embodiment, each DNS reply 256-1 ... 256-R may include state information associated with the corresponding DNS query, while in an alternative embodiment, each DNS reply 256-1 ... 256-R may include a handle to state information associated with the corresponding DNS query. Occasionally, the total size of the corresponding DNS replies may exceed the size of data payload 255 of the Response SuperPacket 240. This overflow may be limited, for example, to a single reply, i.e., the reply associated with the last query contained within Request SuperPacket 220. Rather than sending an additional Response SuperPacket 240 containing only the single reply, the overflow reply may be preferably included in the next Response SuperPacket 240 corresponding to the next Request SuperPacket. Advantageously, header 250 may include appropriate information to determine the extent of the overflow condition. Under peak processing conditions, more than one reply may overflow into the next Response SuperPacket.

[0036] For example, in Response SuperPacket 240, header 250 may include at least two sequence numbers 251-1 and 251-2 and at least two message counts 252-1 and 252-2, grouped as two pairs of complementary fields. While there may be "S" number of sequence number and message count pairs, typically, S is a small number, such as, e.g., 2, 3, 4, etc. Thus, header 250 may include sequence number 251-1 paired with message count 252-1, sequence number 251-2 paired with message count 252-2, etc. Generally, message count 252-1 may reflect the number of replies contained within data payload 255 that are associated with sequence number 251-1. In an embodiment, sequence number 251-1 may be a two-byte field, while message count 252-1 may be a one-byte field.

[0037] In a more specific example, data payload 235 of Request SuperPacket 220 may include seven DNS queries (as depicted in FIG. 2). In one embodiment, sequence number 231-1 may be set to a unique value (e.g., 1024) and message count 232-1 may be set to seven, while sequence number 231-2 and message count 232-2 may be set to zero. In another embodiment, header 230 may contain only one sequence number and one message count, e.g., sequence number 231-1 and message count 232-1 set to 1024 and seven, respectively. Typically, Request SuperPacket 220 may contain all of the queries associated with a particular sequence number.

[0038] Data payload 255 of Response SuperPacket 240 may include seven corresponding DNS replies (as depicted in FIG. 2). In this example, header 250 may include information similar to Request SuperPacket 220, i.e., sequence number 251-1 set to the same unique value (i.e., 1024), message count 252-1 set to seven, and both sequence number 252-2 and message count 252-2 set to zero. However, in another example, data payload 255 of Response SuperPacket 240 may include only five corresponding DNS replies, and message count 252-1 may be set to five instead. The remaining two responses associated with sequence number 1024 may be included within the next Response SuperPacket 240.

[0039] The next Request SuperPacket 240 may include a different sequence number (e.g., 1025) and at least one DNS query, so that the next Response SuperPacket 240 may include the two previous replies associated with the 1024 sequence number, as well as at least one reply associated with the 1025 sequence number. In this example, header 250 of the next Response SuperPacket 240 may include sequence number 251-1 set to 1024, message count 252-1 set to two, sequence number 251-2 set to 1025

and message count 252-2 set to one. Thus, Response SuperPacket 240 may include a total of three replies associated with three queries contained within two different Request SuperPackets.

[0040] FIG. 3 is a detailed block diagram that illustrates a message latency data structure architecture, according to an embodiment of the present invention. Message latency data structure 300 may include information generally associated with the transmission and reception of message 200. In the DNS resolution embodiment, message latency data structure 300 may include latency information about Request SuperPackets and Response SuperPackets; this latency information may be organized in a table format indexed according to sequence number value (e.g., index 301). For example, message latency data structure 300 may include a number of rows N equal to the total number of unique sequence numbers, as illustrated, generally, by table elements 310, 320 and 330. In an embodiment, SuperPacket header sequence numbers may be two bytes in length and define a range of unique sequence numbers from zero to $2^{16}-1$ (i.e., 65,535). In this case, N may be equal to 65,536. Latency information may include Request Timestamp 302, Request Query Count 303, Response Timestamp 304, Response Reply Count 305, and Response Message Count 306. In an alternative embodiment, latency information may also include an Initial Response Timestamp (not shown).

[0041] In an example, table element 320 illustrates latency information for a Request SuperPacket 220 having a single sequence number 231-1 equal to 1024. Request Timestamp 302 may indicate when this particular Request SuperPacket was sent to the LUE. Request Query Count 303 may indicate how many queries were contained within this particular Request SuperPacket. Response Timestamp 304 may indicate when a Response SuperPacket having a sequence number equal to 1024 was received at the PE (e.g., network computer 120-N) and may be updated if more than one Response SuperPacket is received at the PE. Response Reply Count 305 may indicate the total number of replies contained within all of the received Response SuperPackets associated with this sequence number (i.e., 1024). Response Message Count 306 may indicate how many Response SuperPackets having this sequence number (i.e., 1024) arrived at the PE. Replies to the queries contained within this particular Request SuperPacket may be split over several Response SuperPackets, in which case, Response Timestamp 304, Response Reply Count 305, and Response Message Count 306 may be updated as each of the additional Response SuperPackets are received. In

an alternative embodiment, the Initial Response Timestamp may indicate when the first Response SuperPacket containing replies for this sequence number (i.e., 1024) was received at the PE. In this embodiment, Response Timestamp 304 may be updated when additional (i.e., second and subsequent) Response SuperPackets are received.

[0042] Various important latency metrics may be determined from the latency information contained within message latency data structure 300. For example, simple cross-checking between Request Query Count 303 and Response Reply Count 305 for a given index 301 (i.e., sequence number) may indicate a number of missing replies. This difference may indicate the number of queries inexplicably dropped by the LUE. Comparing Request Timestamp 302 and Response Timestamp 304 may indicate how well the particular PE/LUE combination may be performing under the current message load. The difference between the current Request SuperPacket sequence number and the current Response SuperPacket sequence number may be associated with the response performance of the LUE; e.g., the larger the difference, the slower the performance. The Response Message Count 306 may indicate how many Response SuperPackets are being used for each Request SuperPacket, and may be important in DNS resolution traffic analysis. As the latency of the queries and replies travelling between the PEs and LUE increases, the PEs may reduce the number of DNS query packets processed by the system.

[0043] Generally, the LUE may perform a multi-threaded look-up on the incoming, multiplexed Request SuperPackets, and may combine the replies into outgoing, multiplexed Response SuperPackets. For example, the LUE may spawn one search thread, or process, for each active PE and route all the incoming Request SuperPackets from that PE to that search thread. The LUE may spawn a manager thread, or process, to control the association of PEs to search threads, as well as an update thread, or process, to update the database located in memory 104. Each search thread may extract the search queries from the incoming Request SuperPacket, execute the various searches, construct an outgoing Response SuperPacket containing the search replies and send the SuperPacket to the appropriate PE. The update thread may receive updates to the database, from OLTP 140-1, and incorporate the new data into the database. In an alternative embodiment, plurality of network computers 120-1 ... 120-N may send updates to system 100. These updates may be included, for example, within the incoming Request SuperPacket message stream.

[0044] Accordingly, by virtue of the SuperPacket protocol, the LUE may spend less than 15% of its processor capacity on network processing, thereby dramatically increasing search query throughput. In an embodiment, an IBM® 8-way M80 may sustain search rates of 180k to 220k queries per second (qps), while an IBM® 24-way S80 may sustain 400k to 500k qps. Doubling the search rates, i.e., to 500k and 1M qps, respectively, simply requires twice as much hardware, i.e., e.g., two LUEs with their attendant PEs. In another embodiment, a dual Pentium® III 866 MHz multi-processor personal computer operating Red Hat Linux® 6.2 may sustain update rates on the order of 100K/sec. Of course, increases in hardware performance also increase search and update rates associated with embodiments of the present invention, and as manufacturers replace these multiprocessor computers with faster-performing machines, for example, the sustained search and update rates may increase commensurately. Generally, system 100 is not limited to a client or server architecture, and embodiments of the present invention are not limited to any specific combination of hardware and/or software.

[0045] FIG. 4 is a block diagram that illustrates a general database architecture according to an embodiment of the present invention. In this embodiment, database 400 may include at least one table or group of database records 401, and at least one corresponding search index 402 with pointers (indices, direct byte-offsets, etc.) to individual records within the group of database records 401. For example, pointer 405 may reference database record 410.

[0046] In one embodiment, database 400 may include at least one hash table 403 as a search index with pointers (indices, direct byte-offsets, etc.) into the table or group of database records 401. A hash function may map a search key to an integer value which may then be used as an index into hash table 403. Because more than one search key may map to a single integer value, hash buckets may be created using a singly-linked list of hash chain pointers. For example, each entry within hash table 403 may contain a pointer to the first element of a hash bucket, and each element of the hash bucket may contain a hash chain pointer to the next element, or database record, in the linked-list. Advantageously, a hash chain pointer may be required only for those elements, or database records, that reference a subsequent element in the hash bucket.

[0047] Hash table 403 may include an array of 8-byte pointers to individual database records 401. For example, hash pointer 404 within hash table 403 may reference database record 420 as the first element within a hash bucket. Database record 420

may contain a hash chain pointer 424 which may reference the next element, or database record, in the hash bucket. Database record 420 may also include a data length 421 and associated fixed or variable-length data 422. In an embodiment, a null character 423, indicating the termination of data 422, may be included. Additionally, database record 420 may include a data pointer 425 which may reference another database record, either within the group of database records 401 or within a different table or group of database records (not shown), in which additional data may be located.

[0048] System 100 may use various, well-known algorithms to search this data structure architecture for a given search term or key. Generally, database 400 may be searched by multiple search processes, or threads, executing on at least one of the plurality of processors 102-1 ... 102-P. However, modifications to database 400 may not be integrally performed by an update thread (or threads) unless the search thread(s) are prevented from accessing database 400 for the period of time necessary to add, modify, or delete information within database 400. For example, in order to modify database record 430 within database 400, the group of database records 401 may be locked by an update thread to prevent the search threads from accessing database 400 while the update thread is modifying the information within database record 430. There are many well-known mechanisms for locking database 400 to prevent search access, including the use of spin-locks, semaphores, mutexes, etc. Additionally, various off-the-shelf commercial databases provide specific commands to lock all or parts of database 400, e.g., the lock table command in the Oracle 8 Database, manufactured by Oracle Corporation of Redwood Shores, California, etc.

[0049] FIG. 5 is a block diagram that illustrates a general database architecture according to another embodiment of the present invention. In this embodiment, database 500 may include a highly-optimized, read-only, master snapshot file 510 and a growing, look-aside file 520. Master snapshot file 510 may include at least one table or group of database records 511, and at least one corresponding search index 512 with pointers (indices, direct byte-offsets, etc.) to individual records within the group of database records 511. Alternatively, master snapshot file 510 may include at least one hash table 513 as a search index with pointers (indices, direct byte-offsets, etc.) into the table or group of database records 511. Similarly, look-aside file 520 may include at least two tables or groups of database records, including database addition records 521 and database deletion records 531. Corresponding search indices 522 and 532 may be provided, with pointers (indices, direct byte-offsets, etc.) to individual records within the

database addition records 521 and database deletion records 531. Alternatively, look-aside file 520 may include hash tables 523 and 533 as search indices, with pointers (indices, direct byte-offsets, tc.) into database addition records 521 and database deletion records 531, respectively.

[0050] System 100 may use various, well-known algorithms to search this data structure architecture for a given search term or key. In a typical example, look-aside file 520 may include all the recent changes to the data, and may be searched before read-only master snapshot file 510. If the search key is found in look-aside file 520, the response is returned without accessing snapshot file 510, but if the key is not found, then snapshot file 510 may be searched. However, when look-aside file 520 no longer fits in memory 104 with snapshot file 510, search query rates drop dramatically, by a factor of 10 to 50, or more, for example. Consequently, to avoid or minimize any drop in search query rates, snapshot file 510 may be periodically updated, or recreated, by incorporating all of the additions, deletions and modifications contained within look-aside file 520

[0051] Data within snapshot file 510 are not physically altered but logically added, modified or deleted. For example, data within snapshot file 510 may be deleted, or logically "forgotten," by creating a corresponding delete record within database deletion records 531 and writing a pointer to the delete record to the appropriate location in hash table 533. Data within snapshot file 510 may be logically modified by copying a data record from snapshot file 510 to a new data record within database addition records 521, modifying the data within the new entry, and then writing a pointer to the new entry to the appropriate hash table (e.g., hash table 522) or chain pointer within database addition records 521. Similarly, data within snapshot file 510 may be logically added to snapshot file 510 by creating a new data record within database addition records 521 and then writing a pointer to the new entry to the appropriate hash table (e.g., hash table 522) or chain pointer within database addition records 521.

[0052] In the DNS resolution embodiment, for example, snapshot file 510 may include domain name data and name server data, organized as separate data tables, or blocks, with separate search indices (e.g., 511-1, 511-2, 512-1, 512-2, 513-1, 513-2, etc., not shown for clarity). Similarly, look-aside file 520 may include additions and modifications to both the domain name data and the name server data, as well as deletions to both the domain name data and the name server data (e.g., 521-1, 521-2, 522-1, 522-2, 523-1, 523-2, 531-1, 531-2, 532-1, 532-2, 533-1, 533-2, etc., not shown for clarity).

[0053] FIG. 6 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention. Generally, database 600 may be organized into a single, searchable representation of the data. Data set updates may be continuously incorporated into database 600, and deletes or modifications may be physically performed on the relevant database records to free space within memory 104, for example, for subsequent additions or modifications. The single, searchable representation scales extremely well to large data set sizes and high search and update rates, and obviates the need to periodically recreate, propagate and reload snapshot files among multiple search engine computers.

[0054] In a DNS resolution embodiment, for example, database 600 may include domain name data 610 and name server data 620. Domain name data 610 and name server data 620 may include search indices with pointers (indices, direct byte-offsets, etc.) into blocks of variable length records. As discussed above, a hash function may map a search key to an integer value which may then be used as an index into a hash table. Similarly, hash buckets may be created for each hash table index using a singly-linked list of hash chain pointers. Domain name data 610 may include, for example, a hash table 612 as a search index and a block of variable-length domain name records 611. Hash table 612 may include an array of 8-byte pointers to individual domain name records 611, such as, for example, pointer 613 referencing domain name record 620. Variable-length domain name record 620 may include, for example, a next record offset 621, a name length 622, a normalized name 623, a chain pointer 624 (i.e., e.g., pointing to the next record in the hash chain), a number of name servers 625, and a name server pointer 626. The size of both chain pointer 624 and name server pointer 626 may be optimized to reflect the required block size for each particular type of data, e.g., eight bytes for chain pointer 624 and four bytes for name server pointer 626.

[0055] Name server data 630 may include, for example, a hash table 632 as a search index and a block of variable-length name server records 631. Hash table 632 may include an array of 4-byte pointers to individual name server records 631, such as, for example, pointer 633 referencing name server record 640. Variable-length name server record 640 may include, for example, a next record offset 641, a name length 642, a normalized name 643, a chain pointer 644 (i.e., e.g., pointing to the next record in the hash chain), a number of name server network addresses 645, a name server address length 646, and a name server network address 647, which may be, for example, an Internet Protocol (IP) network address. Generally, name server network addresses may

be stored in ASCII (American Standard Code for Information Interchange, e.g., ISO-14962-1997, ANSI-X3.4-1997, etc.) or binary format; in this example, name server network address length 646 indicates that name server network address 647 is stored in binary format (i.e., four bytes). The size of chain pointer 644 may also be optimized to reflect the required name server data block size, e.g., four bytes.

[0056] Generally, both search indices, such as hash tables, and variable-length data records may be structured so that 8-byte pointers are located on 8-byte boundaries in memory. For example, hash table 612 may contain a contiguous array of 8-byte pointers to domain name records 611, and may be stored at a memory address divisible by eight (i.e., an 8-byte boundary, or 8N). Similarly, both search indices, such as hash tables, and variable-length data records may be structured so that 4-byte pointers are located on 4-byte boundaries in memory. For example, hash table 632 may contain a contiguous array of 4-byte pointers to name server records 631, and may be stored at a memory address divisible by four (i.e., a 4-byte boundary, or 4N). Consequently, modifications to database 600 may conclude by updating a pointer to an aligned address in memory using a single uninterruptible operation, including, for example writing a new pointer to the search index, such as a hash table, or writing a new hash chain pointer to a variable-length data record.

[0057] FIG. 7 is a detailed block diagram that illustrates a non-concurrency controlled data structure architecture, according to an embodiment of the present invention. Generally, database 700 may also be organized into a single, searchable representation of the data. Data set updates may be continuously incorporated into database 700, and deletes or modifications may be physically performed on the relevant database records to free space within memory 104, for example, for subsequent additions or modifications. The single, searchable representation scales extremely well to large data set sizes and high search and update rates, and obviates the need to periodically recreate, propagate and reload snapshot files among multiple search engine computers.

[0058] Many different physical data structure organizations are possible. An exemplary organization may use an alternative search index to hash tables for ordered, sequential access to the data records, such as the ternary search tree (trie), or TST, which combines the features of binary search trees and digital search tries. In a text-based applications, such as, for example, whois, domain name resolution using DNS Secure Extensions (Internet Engineering Taskforce Request for Comments: 2535), etc. TSTs

advantageously minimize the number of comparison operations required to be performed, particularly in the case of a search miss, and may yield search performance metrics exceeding search engine implementations with hashing. Additionally, TSTs may also provide advanced text search features, such as, e.g., wildcard searches, which may be useful in text search applications, such as, for example, whois, domain name resolution, Internet content search, etc.

[0059] In an embodiment, a TST may contain a sequence of nodes linked together in a hierarchical relationship. A root node may be located at the top of the tree, related child nodes and links may form branches, and leaf nodes may terminate the end of each branch. Each leaf node may be associated with a particular search key, and each node on the path to the leaf node may contain a single, sequential element of the key. Each node in the tree contains a comparison character, or split value, and three pointers to other successive, or "child," nodes in the tree. These pointers reference child nodes whose split values are less than, equal to, or greater than the node's split value. Searching the TST for a particular key, therefore, involves traversing the tree from the root node to a final leaf node, sequentially comparing each element, or character position, of the key with the split values of the nodes along the path. Additionally, a leaf node may also contain a pointer to a key record, which may, in turn, contain at least one pointer to a terminal data record containing the record data associated with the key (e.g., an IP address). Alternatively, the key record may contain the record data in its entirety. Record data may be stored in binary format, ASCII text format, etc.

[0060] In an embodiment, database 700 may be organized as a TST, including a plurality of fixed-length search nodes 701, a plurality of variable-length key data records 702 and a plurality of variable-length terminal data records 703. Search nodes 701 may include various types of information as described above, including, for example, a comparison character (or value) and position, branch node pointers and a key pointer. The size of the node pointers may generally be determined by the number of nodes, while the size of the key pointers may generally be determined by the size of the variable-length key data set. Key data records 702 may contain key information and terminal data information, including, for example, pointers to terminal data records or embedded record data, while terminal data records 703 may contain record data.

[0061] In an embodiment, each fixed-length search node may be 24 bytes in length. Search node 710, for example, may contain an eight-bit comparison character (or byte

value) 711, a 12-bit character (or byte) position 712, and a 12-bit node type/status (not shown for clarity); these data may be encoded within the first four bytes of the node. The comparison character 711 may be encoded within the first byte of the node as depicted in FIG. 7, or, alternatively, character position 712 may be encoded within the first 12 bits of the node in order to optimize access to character position 712 using a simple shift operation. The next 12 bytes of each search node may contain three 32-bit pointers, i.e., pointer 713, pointer 714 and pointer 715, representing "less than," "equal to," and "greater than" branch node pointers, respectively. These pointers may contain a counter, or node index, rather than a byte-offset or memory address. For fixed-length search nodes, the byte-offset may be calculated from the counter, or index value, and the fixed-length, e.g., counter*length. The final four bytes may contain a 40-bit key pointer 716, which may be a null value indicating that a corresponding key data record does not exist (shown) or a pointer to an existing corresponding key data record (not shown), as well as other data, including, for example, a 12-bit key length and a 12-bit pointer type/status field. Key pointer 716 may contain a byte offset to the appropriate key data record, while the key length may be used to optimize search and insertion when eliminating one-way branching within the TST. The pointer type/status field may contain information used in validity checking and allocation data used in memory management.

[0062] In an embodiment, key data record 750 may include, for example, a variable-length key 753 and at least one terminal data pointer. As depicted in FIG. 7, key data record 750 includes two terminal data pointers: terminal data pointer 757 and terminal data pointer 758. Key data record 750 may be prefixed with a 12-bit key length 751 and a 12-bit terminal pointer count/status 752, and may include padding (not shown for clarity) to align the terminal data pointer 757 and terminal data pointer 758 on an eight-byte boundary in memory 104. Terminal data pointer 757 and terminal data pointer 758 may each contain various data, such as, for example, terminal data type, length, status or data useful in binary record searches. Terminal data pointer 757 and terminal data pointer 758 may be sorted by terminal data type for quicker retrieval of specific resource records (e.g., terminal data record 760 and terminal data record 770). In another embodiment, key data record 740 may include embedded terminal data 746 rather than, or in addition to, terminal data record pointers. For example, key data record 740 may include a key length 741, a terminal pointer count 742, a variable-length key 743, the number of embedded record elements 744, followed by a record element

length 745 (in bytes, for example) and embedded record data 746 (e.g., a string, a byte sequence, etc.) for each of the number of embedded record elements 744.

[0063] In an embodiment, terminal data record 760, for example, may include a 12-bit length 761, a 4-bit status, and a variable-length string 762 (e.g., an IP address). Alternatively, variable length string 762 may be a byte sequence. Terminal data record 760 may include padding to align each terminal data record to an 8-byte boundary in memory 104. Alternatively, terminal data record 760 may include padding to a 4-byte boundary, or, terminal data record 760 may not include any padding. Memory management algorithms may determine, generally, whether terminal data records 760 are padded to 8-byte, 4-byte, or 0-byte boundaries. Similarly, terminal data record 770 may include a 12-bit length 771, a 4-bit status, and a variable-length string 772 (e.g., an IP address).

[0064] Generally, both search indices, such as TSTs, and data records may be structured so that 8-byte pointers are located on 8-byte boundaries in memory. For example, key pointer 726 may contain an 8-byte (or less) pointer to key data record 740, and may be stored at a memory address divisible by eight (i.e., an 8-byte boundary, or 8N). Similarly, both search indices, such as TSTs, and data records may be structured so that 4-byte pointers are located on 4-byte boundaries in memory. For example, node branch pointer 724 may contain a 4-byte (or less) pointer to node 730, and may be stored at a memory address divisible by four (i.e., a 4-byte boundary, or 4N). Consequently, modifications to database 700 may conclude by updating a pointer to an aligned address in memory using a single uninterruptible operation, including, for example writing a new pointer to the search index, such as a TST node, or writing a new pointer to a data record.

[0065] FIG. 8 is a detailed block diagram that illustrates another data structure architecture, according to an embodiment of the present invention. As above, database 800 may also be organized into a single, searchable representation of the data. Data set updates may be continuously incorporated into database 800, and deletes or modifications may be physically performed on the relevant database records to free space within memory 104, for example, for subsequent additions or modifications. The single, searchable representation scales extremely well to large data set sizes and high search and update rates, and obviates the need to periodically recreate, propagate and reload snapshot files among multiple search engine computers.

[0066] Other search index structures are possible for accessing record data. In an embodiment, database 800 may use an alternative ordered search index, organized as an ordered access key tree (i.e., "OAK tree"). Database 800 may include, for example, a plurality of variable-length search nodes 801, a plurality of variable-length key records 802 and a plurality of variable-length terminal data records 803. Search nodes 801 may include various types of information as described above, such as, for example, search keys, pointers to other search nodes, pointers to key records, etc. In an embodiment, plurality of search nodes 801 may include vertical and horizontal nodes containing fragments of search keys (e.g., strings), as well as pointers to other search nodes or key records. Vertical nodes may include, for example, at least one search key, or character, pointers to horizontal nodes within the plurality of search nodes 801, pointers to key records within the plurality of key records 802, etc. Horizontal nodes may include, for example, at least two search keys, or characters, pointers to vertical nodes within the plurality of search nodes 801, pointers to horizontal nodes within the plurality of search nodes 801, pointers to key records within the plurality of key records 802, etc. Generally, vertical nodes may include a sequence of keys (e.g., characters) representing a search key fragment (e.g., string), while horizontal nodes may include various keys (e.g., characters) that may exist at a particular position within the search key fragment (e.g., string).

[0067] In an embodiment, plurality of search nodes 801 may include vertical node 810, vertical node 820 and horizontal node 830. Vertical node 810 may include, for example, a 2-bit node type 811 (e.g., "10"), a 38-bit address 812, an 8-bit length 813 (e.g., "8"), an 8-bit first character 814 (e.g., "I") and an 8-bit second character 815 (e.g., "null"). In this example, address 812 may point to the next node in the search tree, i.e., vertical node 820. In an embodiment, 38-bit address 812 may include a 1-bit terminal/nodal indicator and a 37-bit offset address to reference one of the 8-byte words within a 1 Tbyte ($\sim 10^{12}$ byte) address space of memory 104. Accordingly, vertical node 810 may be eight bytes (64 bits) in length, and, advantageously, may be located on an 8-byte word boundary within memory 104. Generally, each vertical node within plurality of search nodes 801 may be located on an 8-byte word boundary within memory 104.

[0068] A vertical node may include a multi-character, search key fragment (e.g., string). Generally, search keys without associated key data records may be collapsed into a single vertical node to effectively reduce the number of vertical nodes required within plurality of search nodes 801. In an embodiment, vertical node 810 may include eight

bits for each additional character, above two characters, within the search key fragment, such as, for example, 8-bit characters 816-1, 816-2 ... 816-N (shown in phantom outline). Advantageously, vertical node 810 may be padded to a 64-bit boundary within memory 104 in accordance with the number of additional characters located within the string fragment. For example, if nine characters are to be included within vertical node 810, then characters one and two may be assigned to first character 814 and second character 815, respectively, and 56 bits of additional character information, corresponding to characters three through nine, may be appended to vertical node 810. An additional eight bits of padding may be included to align the additional character information on an 8-byte word boundary.

[0069] Similarly, vertical node 820 may include, for example, a 2-bit node type 821 (e.g., "10"), a 38-bit address 822, an 8-bit length 823 (e.g., "8"), an 8-bit first character 824 (e.g., "a") and an 8-bit second character 825 (e.g., "null"). In this example, address 822 may point to the next node in the search tree, i.e., horizontal node 830. Accordingly, vertical node 820 may be eight bytes in length, and, advantageously, may be located on an 8-byte word boundary within memory 104. Of course, additional information may also be included within vertical node 820 if required, as described above with reference to vertical node 810.

[0070] Horizontal node 830 may include, for example, a 2-bit node type 831 (e.g., "01"), a 38-bit first address 832, an 8-bit address count 833 (e.g., 2), an 8-bit first character 834 (e.g., "."), an 8-bit last character 835 (e.g., "w"), a variable-length bitmap 836 and a 38-bit second address 837. In this example, first character 834 may include a single character, "." representing the search key fragment "la" defined by vertical nodes 810 and 820, while last character 831 may include a single character "w," representing the search key fragment "law" defined by vertical nodes 810 and 820, and the last character 835 of horizontal node 830. First address 832 may point to key data record 840, associated with the search key fragment "la," while second address 837 may point to key data record 850 associated with the search key fragment "law."

[0071] Bitmap 836 may advantageously indicate which keys (e.g., characters) are referenced by horizontal node 830. A "1" within a bit position in bitmap 836 indicates that the key, or character, is referenced by horizontal node 830, while a "0" within a bit position in bitmap 836 may indicate that the key, or character, is not referenced by horizontal node 830. Generally, the length of bitmap 836 may depend upon the number

of sequential keys, or characters, between first character 834 and last character 835, inclusive of these boundary characters. For example, if first character 834 is "a" and last character 835 is "z," then bitmap 836 may be 26 bits in length, where each bit corresponds to one of the characters between, and including, "a" through "z." In this example, additional 38-bit addresses would be appended to the end of horizontal node 830, corresponding to each of the characters represented within bitmap 836. Each of these 38-bit addresses, as well as bitmap 836, may be padded to align each quantity on an 8-byte word boundary within memory 104. In an embodiment, the eight-bit ASCII character set may be used as the search key space so that bitmap 836 may be as long as 256 bits (i.e., 2^8 bits or 32 bytes). In the example depicted in FIG. 8, due to the special reference character "*" and address count 833 of "2," bitmap 836 may be two bits in length and may include a "1" in each bit position corresponding to last character 835.

[0072] In an embodiment, and as discussed with reference to key data record 750 (FIG. 7), key data record 850 may include, for example, a variable-length key 853 and at least one terminal data pointer. As depicted in FIG. 8, key data record 850 includes two terminal data pointers, terminal data pointer 857 and terminal data pointer 858. Key data record 850 may be prefixed with a 12-bit key length 851 and a 12-bit terminal pointer count/status 852, and may include padding (not shown for clarity) to align the terminal data pointer 857 and terminal data pointer 858 on an 8-byte boundary in memory 104. Terminal data pointer 857 and terminal data pointer 858 may each contain a 10-bit terminal data type and other data, such as, for example, length, status or data useful in binary record searches. Terminal data pointer 857 and terminal data pointer 858 may be sorted by terminal data type for quicker retrieval of specific resource records (e.g., terminal data record 860 and terminal data record 870).

[0073] In another embodiment, and as discussed with reference to key data record 740 (FIG. 7), key data record 840 may include embedded terminal data 846 rather than a terminal data record pointer. For example, key data record 840 may include a key length 841, a terminal pointer count 842, a variable-length key 843, the number of embedded record elements 844, followed by a record element length 845 (in bytes, for example) and embedded record data 846 (e.g., a string, a byte sequence, etc.) for each of the number of embedded record elements 844.

[0074] In another embodiment, and as discussed with reference to terminal data record 760 (FIG. 7), terminal data record 860, for example, may include a 12-bit length 861, a

4-bit status, and a variable-length string 862 (e.g., an IP address). Alternatively, variable length string 862 may be a byte sequence. Terminal data record 860 may include padding (not shown for clarity) to align each terminal data record to an 8-byte boundary in memory 104. Alternatively, terminal data record 860 may include padding (not shown for clarity) to a 4-byte boundary, or, terminal data record 860 may not include any padding. Memory management algorithms may determine, generally, whether terminal data records 760 are padded to 8-byte, 4-byte, or 0-byte boundaries. Similarly, terminal data record 870 may include a 12-bit length 871, a 4-bit status, and a variable-length string 872 (e.g., an IP address).

[0075] Generally, both search indices, such as OAK trees, and data records may be structured so that 8-byte pointers are located on 8-byte boundaries in memory. For example, vertical node 810 may contain an 8-byte (or less) pointer to vertical node 820, and may be stored at a memory address divisible by eight (i.e., an 8-byte boundary, or 8N). Similarly, both search indices, such as OAK trees, and data records may be structured so that 4-byte pointers are located on 4-byte boundaries in memory. Consequently, modifications to database 800 may conclude by updating a pointer to an aligned address in memory using a single uninterruptible operation, including, for example writing a new pointer to the search index, such as an OAK trees node, or writing a new pointer to a data record.

[0076] The various embodiments discussed above with reference to FIG. 8 present many advantages. For example, an OAK tree data structure is extremely space efficient and 8-bit clean. Regular expression searches may be used to search vertical nodes containing multi-character string fragments, since the 8-bit first character (e.g., first character 814), the 8-bit second character (e.g., second character 8-15) and any additional 8-bit characters (e.g., additional characters 816-1 ... 816-N) may be contiguously located within the vertical node (e.g., vertical node 810). Search misses may be discovered quickly, and, no more than N nodes may need to be traversed to search for an N-character length search string.

[0077] FIG. 9 is a top level flow diagram that illustrates a method for processing query messages received over a network, according to an embodiment of the present invention.

[0078] A plurality of queries may be extracted (900) from a plurality of query messages received from a plurality of users over the network. In the DNS resolution embodiment, for example, one of the plurality of network computers 120-1 ... 120-N (e.g., network computer 120-1) may extract (900) a plurality of DNS queries from a plurality of DNS query messages received from a plurality of users over wide area network 124. For example, network computer 120-1 may accumulate and process DNS query messages to identify valid DNS queries, as well as to reject invalid DNS queries. In one embodiment, network computer 120-1 may simply ignore an invalid DNS query, while in another embodiment, network computer 120-1 may return an error message to the appropriate user over wide area network 124. Of course, corrupted network packets containing these DNS queries may also be detected and the appropriate action may be taken based on the underlying network protocols operating within wide area network 124.

[0079] A number of queries included in the plurality of queries may be determined (910). In the DNS resolution embodiment, for example, network computer 120-1 may accumulate valid DNS queries and determine (910) that the number of accumulated DNS queries equals a maximum number of queries. For example, in one embodiment, the maximum number of queries may be 30, while in another embodiment, the maximum number of queries may be 60 or higher. Generally, the maximum number of queries depends upon the average size of the DNS reply, as well as the MTU of local area network 122, as discussed above.

[0080] In an alternative embodiment, network computer 120-1 may accumulate valid DNS queries over a predetermined time period and then determine (910) the number of queries. For example, the predetermined time period may be 2 ms. In this embodiment, the number of valid DNS queries, accumulated over the predetermined time period, may be less than the maximum number of queries discussed above. In another embodiment, a combination of a predetermined time period and a maximum number of queries may be used in order to optimize the number of accumulated queries without incurring a significant time penalty during accumulation.

[0081] A current sequence number may be associated (920) with the plurality of queries. In the DNS resolution embodiment, for example, network computer 120-1 may associate (920) a current sequence number with the accumulated DNS queries. In an embodiment, the sequence number may be represented by a 16-bit unsigned integer, so

that 2^{16} (i.e., 65,536) different sequence numbers may be defined. For example, if the sequence number associated with the previous set of accumulated queries was "1024," then the sequence number associated with the current set of accumulated queries may be "1025." Generally, each set of accumulated queries may be assigned a different sequence number, which may be incremented by a predetermined amount (e.g., 1) for each successive accumulated query set. For an unsigned, 16-bit integer sequence number, once the maximum sequence number is reached (i.e., 65,535), the sequence number may be reset to 0.

[0082] A request message may be created (930) including the plurality of queries, a first sequence number and a first message count number. In the DNS resolution embodiment, for example, network computer 120-1 may create (930) a Request SuperPacket 220, including a sequence number 231-1 (e.g., 1024) and a message count number 232-1 (e.g., 7). Sequence number 231-2 and message count 232-2 may also be included within Request SuperPacket 220 and may each be set to 0. In this example, seven DNS queries may be included within data payload 235 (e.g., DNS queries 236-1 ... 236-Q, where Q equals 7). Generally, the DNS queries may include state information, such as, for example, the source address, the protocol type, etc.

[0083] The request message may be sent (940) to a search engine. In the DNS resolution embodiment, for example, network computer 120-1 (e.g., a PE) may send (940) the Request SuperPacket 220 to system 100 (e.g., an LUE) over local area network 122 for processing.

[0084] A response message may be received (950) from the search engine, including a plurality of replies, a second sequence number, a second message count, a third sequence number and third message count. In the DNS resolution embodiment, for example, network computer 120-1 may receive (950) a Response SuperPacket 230 from system 100 over local area network 122. The Response SuperPacket 230 may include, for example, a sequence number 251-1 (e.g., 1024), a message count number 252-1 (e.g., 7), a sequence number 251-2 (e.g., 0), a message count number 252-2 (e.g., 0) and data payload 255 containing DNS replies corresponding to the DNS queries contained within the Request SuperPacket 220. In this example, seven DNS replies may be included within data payload 255 (e.g., DNS replies 256-1 ... 256-R, where R equals 7). Similarly, the DNS replies may include state information provided by the DNS queries, such as, for example, the source address, the protocol type, etc.

[0085] In another example, if the total size of the corresponding DNS replies exceeds the maximum length of data payload 255, then a lesser number of DNS replies may be included within data payload 255. In this example, only five DNS replies may be contained within data payload 255 of Response SuperPacket 230. Message count number 252-1 may be set to five, and five DNS replies, corresponding to five DNS queries, may be included within data payload 255. The remaining two DNS replies may be included in a subsequent Response SuperPacket 230 sent from system 100 to network computer 120-1. Thus, in this example, network computer 120-1 may receive (950) a Response SuperPacket 230 containing five DNS replies corresponding to five DNS queries sent within Request SuperPacket 220.

[0086] Importantly, the subsequent Response SuperPacket 230 may contain not only the remaining two DNS replies associated with the sequence number "1024," but also additional DNS replies corresponding to a set of subsequent DNS queries associated with the sequence number "1025." For example, a subsequent Request SuperPacket 220 may contain a sequence number 231-1 set to 1025, a message count number 232-2 set to 3, and three DNS queries accumulated over a subsequent predetermined time period. In this example, the subsequent Response SuperPacket 230 may contain sequence number 251-1 set to 1024, message count number 252-1 set to 2, and the remaining two DNS replies associated with sequence number "1024." Additionally, Response SuperPacket 230 may contain sequence number 251-2 set to 1025, message count number 252-2 set to 3, and three DNS replies associated with sequence number "1025." Thus, subsequent Response SuperPacket 230 may include a total of five DNS replies associated with sequence numbers "1024" and "1025."

[0087] A plurality of reply messages may be created (960) from the plurality of replies and sent (970) to the plurality of users. In the DNS resolution embodiment, for example, network computer 120-1 may create (960) a DNS reply message from each of the DNS replies contained within data payload 255 of Response SuperPacket 230, and send (970) the DNS reply messages to the appropriate users over wide area network 124. The DNS replies may include state information, as noted above, to facilitate the creation of the plurality of reply messages.

[0088] In a further embodiment, a message latency associated with a sequence number may be determined (980). In the DNS resolution embodiment, for example, network computer 120-1 may determine (980) a message latency associated with a sequence

number (e.g., 1024). Network computer 120-1 may include message latency data structure 300 for that purpose. For clarity, the top level flow diagram illustrated in FIG. 9 is extended to FIG. 10 though flow diagram connection symbol "A."

[0089] FIG. 10 is a top level flow diagram that illustrates a method for determining a message latency associated with a sequence number, according to an embodiment of the present invention.

[0090] A request timestamp may be updated (1000) based on the request message. In the DNS resolution embodiment, for example, a Request SuperPacket 220 may include a sequence number 231-1 (e.g., 1024) and a message count number 232-1 (e.g., 7). Sequence number 231-2 and message count 232-2 may be included within Request SuperPacket 220 and set to 0, and seven DNS queries may be included within data payload 235 (e.g., DNS queries 236-1 ... 236-Q, where Q equals 7). Table element 320 includes message latency information for sequence number 1024 (e.g., index 321 equals 1024). In this example, network computer 120-1 may update (1000) request timestamp 322 when Request SuperPacket 220 is sent to system 100 over local area network 122.

[0091] A response timestamp may be updated (1010) based on the response message. In the DNS resolution embodiment, for example, a Response SuperPacket 240 may include a sequence number 251-1 (e.g., 1024) and a message count number 252-1 (e.g., 7). Sequence number 251-2 and message count 252-2 may also be included within Response SuperPacket 240 and may each be set to 0, and seven DNS replies may be included within data payload 255 (e.g., DNS replies 256-1 ... 256-R, where R equals 7). In this example, network computer 120-1 may update (1010) response timestamp 324 when Response SuperPacket 220 is received from system 100 over local area network 122.

[0092] The request timestamp and the response timestamp may be compared (1020) to determine the message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, network computer 120-1 may compare request timestamp 322 and response timestamp 324, for sequence number 1024, to determine the message latency, or time delay, between these two messages. The time delay may advantageously indicate how well network computer 120-1 (e.g., a PE) and system 100 (e.g., an LUE) may be performing under the current DNS query message

load arriving over wide area network 124. A small time delay, such as, for example, 250 ms, may indicate an acceptable message latency. A large time delay, such as, for example, a time difference greater than 250 ms, may indicate below-optimal performance and an excessive message latency.

[0093] In another embodiment, a query count may be updated (1030) based on the request message, a reply count may be updated (1040) based on the response message, and the query count may be compared (1050) to the reply count to determine the message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, Request SuperPacket 220 may include sequence number 231-1 equal to 1024, message count 232-1 equal to 7, and seven DNS queries 236-1 ... 236-Q (where Q equals 7), while Response SuperPacket 240 may include sequence number 251-1 equal to 1024, message count 252-1 equal to 7, and seven DNS replies 256-1 ... 256-R (where R equals 7). Network computer 120-1 may update (1030) query count 323 (e.g., 7) when Request SuperPacket 220 is sent to system 100 over local area network 122. Similarly, network computer 120-1 may update (1040) reply count 325 (e.g., 7) when Response SuperPacket 240 is received from system 100 over local area network 122. Network computer 120-1 may then compare (1050) query count 323 and reply count 325 to determine whether any replies are missing from Response SuperPacket 240. In this example, there are seven queries and seven corresponding replies. Consequently, no replies have been inexplicably dropped by system 100 and the message latency may be minimal.

[0094] In another embodiment, a response count may be updated (1060) based on the response message and the response count may be compared (1070) to a predetermined response count to determine the message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, Request SuperPacket 220 may include sequence number 231-1 equal to 1024, message count 232-1 equal to 7, and seven DNS queries 236-1 ... 236-Q (where Q equals 7), while Response SuperPacket 240 may include sequence number 251-1 equal to 1024, message count 252-1 equal to 7, and seven DNS replies 256-1 ... 256-R (where R equals 7). In this particular example, Response SuperPacket 240 contains replies associated with a single sequence number (i.e., 1024).

[0095] In this example, network computer 120-1 may update (1060) response count 326 when a Response SuperPacket 240 containing a sequence number equal to 1024 is

received from system 100 over wide area network 124. Accordingly, network computer 120-1 may update (1060) response count 326, from an initial value of "0" to a value of "1," when the first Response SuperPacket 240 containing a sequence number equal to 1024 is received from system 100. Network computer 120-1 may update (1060) response count 326 as each subsequent Response SuperPacket 240, containing a sequence number equal to 1024, is received. Network computer 120-1 may then compare (1070) response count 326 to a predetermined response count (e.g., 1) to determine the message latency associated with sequence number 1024.

[0096] Generally, the predetermined response count may depend upon the anticipated network traffic transmitted within local area network 122, and may indicate the expected number of Response SuperPackets that are required for each Request SuperPacket. In this particular example, there are seven queries within a single Request SuperPacket 230 and seven corresponding replies contained within a single Response SuperPacket 240. Consequently, response count 326 equals the predetermined response count (i.e., 1), which may indicate a minimal message latency.

[0097] FIG. 11 is a top level flow diagram that illustrates a method for determining a message latency associated with a sequence number, according to an embodiment of the present invention. For clarity, the top level flow diagram illustrated in FIG. 10 is extended to FIG. 11 through flow diagram connection symbol "B."

[0098] An additional response message may be received (1100) from the search engine, including an additional plurality of replies, a fourth sequence number equal to the first sequence number and a fourth message count greater than zero. In the DNS resolution embodiment, for example, two or more Response SuperPackets may be used to transfer all of the replies corresponding to the queries contained within a single Request SuperPacket. For example, a first Request SuperPacket 220 may include sequence number 231-1 equal to 1024, message count 232-1 equal to 7, and seven DNS queries 236-1 ... 236-Q (where Q equals 7), while a second Request SuperPacket 220 may include sequence number 231-1 equal to 1025, message count 232-1 equal to 3, and three DNS queries 236-1 ... 236-Q (where Q equals 3). A first Response SuperPacket 240 may include sequence number 251-1 equal to 1024, message count 252-1 equal to 5, and five DNS replies 256-1 ... 256-R (where R equals 5), while a second Response SuperPacket 240 may include sequence number 251-1 equal to 1024, message count 252-1 equal to 2, sequence number 251-2 equal to 1025,

message count 252-2 equal to 3, and five DNS replies 256-1 ... 256-R (where R equals 5), the first two replies being associated with sequence number "1024" and the latter three replies being associated with sequence number "1025." Thus, in this example, the DNS replies corresponding to the DNS queries contained within the first Request SuperPacket 220 may be divided between two Response SuperPackets 240.

[0099] In one embodiment, a response timestamp may be updated (1110) based on the additional response message, and the request timestamp and the updated response timestamp may be compared (1120) to determine an updated message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, network computer 120-1 may update (1110) response timestamp 324 when the second Response SuperPacket 240, containing two replies associated with sequence number "1024," is received from system 100 over local area network 122. Network computer 120-1 may compare (1120) updated response timestamp 324 and request timestamp 322 to determine an updated message latency, or updated time delay, associated with sequence number "1024." As noted above, the time delay may generally indicate how well network computer 120-1 (e.g., a PE) and system 100 (e.g., an LUE) may be performing under the current query message load arriving over wide area network 124. A large time delay, such as, for example, a time difference greater than 250 ms, may indicate below-optimal performance. In this example, request timestamp 322 may equal 01:00.500 and updated response timestamp 324 may equal 01:01.750, so that an updated time delay of 1.250 seconds may result, which may indicate an excessive message latency.

[00100] In another embodiment, a reply count may be updated (1130) based on the additional response message, and the updated reply count may be compared (1140) to the query count to determine an updated message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, network computer 120-1 may update (1130) reply count 325 when the second Response SuperPacket 240, containing two replies associated with sequence number 1024, is received from system 100 over local area network 122. Network computer 120-1 may compare (1140) updated reply count 325 and query count 323 to determine whether any replies are missing from the additional Response SuperPacket 240. In this example, there are seven queries sent within Request SuperPacket 220, five corresponding replies within the first Response SuperPacket 240 and two corresponding replies within the second

Response SuperPacket 240. Consequently, in this example, no replies have been dropped by system 100 and the message latency may be minimal.

[00101] In another embodiment, a response count may be updated (1150) based on the response message, and the updated response count may be compared (1160) to a predetermined response count to determine an updated message latency associated with a particular sequence number. In the DNS resolution embodiment, for example, network computer 120-1 may update (1150) response count 326 when the second Response SuperPacket 240, containing two replies associated with sequence number "1024," is received from system 100 over local area network 122. Network computer 120-1 may compare (1140) updated response count 326 to the predetermined response count (e.g., 1) to determine the message latency associated with sequence number 1024. As noted above, the predetermined response count may generally depend upon the anticipated network traffic transmitted within local area network 122, and may indicate the expected number of Response SuperPackets 240 that are required for each Request SuperPacket 220. In this example, there are seven queries within the Request SuperPacket 220, five corresponding replies within the first Response SuperPacket 240 and two corresponding replies within the second Response SuperPacket 240. In this case, response count 326 is greater than the predetermined response count (i.e., 1), and may indicate a marginal message latency.

[00102] Several embodiments of the present invention are specifically illustrated and described herein. However, it will be appreciated that modifications and variations of the present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.